Satisficing, maximizing and bounded rationality

Abstract

It is often argued that maximizing normative theories are inappropriate for bounded agents. In this paper, I argue that traditional motivations for theories of bounded rationality are substantially compatible with some traditional maximizing normative theories. This finding expands the audience for theories of bounded rationality by making room for advocates of maximizing normative theories, and has implications for the explanatory status of Bayesian cognitive science, the relationship between coherence and bounded rationality, and the nature of satisficing.

1 Introduction

Bounded agents satisfice. They do not maximize. The first sentence expresses a truism. The second is often treated as such.

> Information-processing systems typically need to *satisfice* rather than optimize. (Gigerenzer and Goldstein 1996, p. 651)

> When we leave the thin theory of rationality, the link between rationality and optimality is completely broken. (Elster 1983, p. 18)

> Human rationality [is] non-maximizing and non-optimizing, but rather bounded, heuristic and satisficing. (Viale 2021, p. 1)

> Ecological rationality, a development of Herbert Simon's . . . notion of bounded rationality, provides a fully-fledged alternative to the view that rational decisions under environmental uncertainty are optimal probabilistic decisions. (Brighton 2020, p. 268)

> Bounded rationality dispenses with the notion of optimization. (Gigerenzer and Selten 2001b, p. 6)

My goal in this paper is to show that contemporary knowledge about bounded rationality is surprisingly compatible with a traditional philosophical form of maximization.

In doing this, my aim is not to criticize non-maximizing approaches to bounded rationality or to force rivals into maximizing theories. My motivation is as follows. A

large number of philosophers find themselves attracted to maximizing normative theories. These philosophers may be led to think that insights from bounded rationality are not compatible with their maximizing normative views. By showing how traditional maximizing approaches can be made compatible with contemporary insights from the bounded tradition, I hope to show how these insights can be made accessible to philosophers committed to maximizing normative theories.

Here is the plan. Section 2 makes five remarks to clarify my project and the relevant notion of maximization. Sections 3-6 discuss leading reasons why bounded rationality has been taken to be incompatible with maximization. Section 7 concludes by drawing out philosophical implications of this discussion: an alternative to traditional approaches such as minimal rationality (Cherniak 1981, 1986; Hoek forthcoming) and satisficing consequentialism (Byron 2004; Chappell 2019; Slote 1984) (Section 7.1), a transformed understanding of the possibility and contents of Bayesian approaches to bounded rationality (Section 7.2), a new project of determining how many extant maximizing approaches can be made consistent with insights from bounded rationality (Section 7.3), and a need to further clarify and distinguish senses of satisficing (Section 7.4).

2 Clarifications

Before beginning, I want to make five clarificatory remarks. First, it is important to distinguish the descriptive project of saying how bounded agents in fact cognize from the normative project of assessing their cognition (Bermúdez 2009). My project is normative, rather than descriptive. However, many bounded rationality theorists (Gigerenzer and Selten 2001b; Sturm 2012; Thorstad 2024) share with Bayesian rational analysts (Anderson 1990; Chater and Oaksford 1999; Icard 2018) the assumption that human cognition is often fairly rational. If that is right, then normative theorizing can often provide a good defeasible starting point for descriptive theorizing. Nevertheless, it remains important to clearly distinguish between normative and descriptive projects.

Second, for concreteness it will help to fix a particular maximizing normative standard. In this paper, I am interested in the prospects for a traditional form of maximizing consequentialism. Say that option X is maximal for S at t if no alternative available to S at t is better than X. Maximizing consequentialism holds:

> **Maximizing Consequentialism:** For all agents S, acts X and times t:
>
> > **Maximizing Criterion of Correctness:** S may (must) X at t just in case X is maximal (uniquely maximal) for S at t.
> >
> > **Information-Sensitive Value:** The value of S's X-ing at t is the expected value of X's consequences, where the expectation in question is taken relative to evidential probabilities. That is, $V(X) = \sum_w Pr_E(w)V(X,w)$, where Pr_E are the evidential probabilities generated by S's total evidence at t.[1]

I suspect that the discussion in this paper can generalize to many, but perhaps not all leading maximizing views. Section 7.3 discusses the prospects for some other maximizing views.

Third, my aim is only to show that maximizing consequentialism is compatible with existing knowledge about bounded rationality, not that maximizing consequentialism is required to make sense of bounded rationality. For this reason, I will not make novel arguments in favor of maximizing consequentialism or against competing views. My goal is to expand the range of normative theories which can be used to recover traditional insights about bounded rationality, not to shrink it.

Fourth, some elements of my project may be partly anticipated by earlier work. Many theorists including Herbert Simon have admitted that the line between maximizing and satisficing approaches is less sharp than we might suppose.[2] Some have gone further and

[1]For discussions of information-sensitivity see Charlow (2013); Kolodny and MacFarlane (2010) and Yalcin (2016) among others. Readers who prefer objective or subjective readings of deontic modals are invited to explore the implications of these readings throughout the paper.

[2]For example, Simon (1978, p. 8) sets out "to show how large a plot of common ground is shared by

discussed the possibility of taking maximizing theories on board as normative standards, which is the proposal that I endorse.[3] I think that this background helps to reduce the seeming heterodoxy of my approach, without eliminating the need for a new and detailed defense. We saw in Section 1 that some elements of my proposal are deeply controversial, and we will see in Section 7 that this approach has important philosophical implications that have not been sufficiently emphasized in existing philosophical and scientific debates. For this reason, it is important to present a careful defense of the view and to take time exploring its implications.

Fifth, because debates about bounded rationality are highly interdisciplinary, some of the argumentative moves made in this paper may seem familiar or obvious to philosophers. This does not remove the necessity of making those moves, since we will see that all of these moves have been missed by key texts in the bounded tradition, leading to the proliferation of normative challenges to maximizing consequentialism which are not decisive. One productive way to understand some of the more familiar elements of my discussion would be as a renewed demonstration of the importance of interdisciplinary dialogue and the ability of normative philosophy to make important contributions to scientific debates.

Summing up, my project is the normative project of saying how bounded agents should cognize. I concentrate on the prospects for a traditional form of maximizing consequentialism. My aim is to show that maximizing consequentialism is compatible with considerations of bounded rationality thought to rule out maximizing approaches. This project builds on earlier scientific and philosophical work, while providing novel elaboration, defense and philosophical implications. Let us now consider how maximizing consequentialism fares against traditional objections by bounded rationality theorists.

---

optimizing and satisficing analyses." See also Samuels et al. (2002) on common ground between competing approaches.

[3]For example, Thomas Sturm (2012, pp. 77-8) notes that "formal rules, such as Bayes' theorem or principles of optimization (such as maximizing expected utility) continue to play an important role even within attempts to construct fast and frugal heuristics, namely as a normative standard." However, Sturm also holds that such norms play different functional roles from traditional norms of bounded rationality.

3 Modest changes

At least two familiar complaints against maximizing theories can be handled through at most modest changes in the view.

3.1 All options are dominated

In some cases, there is no best option for agents to take. Consider:

> **(Everbetter Wine)** An immortal possesses a bottle of wine which gets better every year. When should she drink it? (Pollock 1983)

> **(Finite Bliss)** God asks you to name a number n. You will be rewarded with n years of bliss, followed by annihilation. Which number should you name? (Tucker 2016)

For any n, the option of drinking the wine in year n or asking for n years of bliss is suboptimal, because it would be better to opt for $n + 1$ years of bliss or wine maturation. But this means that no matter what the agent chooses, her choice will be worse than another option she could have selected. This means that maximizing consequentialism rules out all possible options in both cases. Does this pose a decisive objection to maximizing consequentialism (Elster 1983; Schmidtz 2004)?

We can react to such cases in two ways. On the one hand, some theorists have recommended holding our ground, retaining maximizing consequentialism and treating the cases as rational dilemmas in which no option is permissible (Sorensen 2006). This would require no change to maximizing consequentialism.

One way to defend this stalwart reaction is to note that the cases in question threaten not only maximizing consequentialism, but also much weaker principles. For example, consider:

> **(Very Weak Statewise Dominance)** For any agent S and options o, o' available

to S, if o is guaranteed to provide a strictly better outcome than o' in every state of the world, then o' is impermissible.

The intuition behind Very Weak Statewise Dominance is that it would be impermissible to choose o' because it would be better to choose o instead. Dominance principles are widely regarded as among the most minimal and normatively compelling decision-theoretic principles, and Very Weak Statewise Dominance is one of the weakest dominance principles on offer. But note that all options in Finite Bliss and Everbetter Wine are dominated, so that if we do not treat these cases as rational dilemmas, we need not only to reject maximizing consequentialism, but also a great number of decision-theoretic principles which include or imply Very Weak Statewise Dominance. And that is a very costly thing to do.

On the other hand, we might take cases such as Finite Bliss and Everbetter Wine to show that maximizing consequentialism is inapplicable in cases where there is no best option. That would involve a gentle relaxation of maximizing consequentialism, but it would not be a devastating change, since cases like Everbetter Wine and Finite Bliss are rarely encountered in our daily lives. Nor would it be an unprincipled move, for when asked why we decided to relax maximizing consequentialism in exactly these cases, we would reply that we relaxed maximizing consequentialism because it cannot be coherently applied to these cases.[4] At the same time, it is important to recall that relaxing maximizing consequentialism comes with a number of costs. In this case, we would also have to deny Very Weak Statewise Dominance, and that is something we may not be willing to do.

3.2 Incommensurability

Some theorists object that maximizing theories do not allow for the possibility of incommensurable options (Gigerenzer and Selten 2001; Klein 2001; Schmidtz 1995). Here the problem cannot be with the Maximizing Criterion of Correctness:

Maximizing Criterion of Correctness: S may (must) X at t just in case X is

[4] Alan Hajek (2014) defends a similar position.

maximal (uniquely maximal) for S at t.

The Maximizing Criterion of Correctness allows for incommensurable options by allowing multiple options to be maximal. The Maximizing Criterion of Correctness would put pressure against incommensurability if it were not formulated in terms of maximality:

> **Restrictive Maximizing Criterion of Correctness:** S may (must) X at t just in case X is at least as good as (better than) all alternatives to X for S at t.

The Restrictive Maximizing Criterion of Correctness treats cases of incommensurability as rational dilemmas, since incommensurable maximal options will not be better than, or at least as good as one another. That may be unappealing. It is precisely to block this consequence that many theorists prefer the Maximizing Criterion of Correctness to the Restrictive Maximizing Criterion of Correctness.

We are not out of the woods yet, for incommensurability does conflict with my statement of information-sensitivity:

> **Information-Sensitive Value:** The value of S's X-ing at t is the expected value of X's consequences, where the expectation in question is taken relative to evidential probabilities. That is, $V(X) = \sum_w Pr_E(w)V(X, w)$, where Pr_E are the evidential probabilities generated by S's total evidence at t.

However, incommensurability is not incompatible with the core insight of information-sensitivity:

> **Value is Information-Sensitive:** The value of S's X-ing at t is assessed against S's total evidence.

The core claim that Value is Information-Sensitive is compatible with incommensurability, but it will have to be given a different decision-theoretic articulation.

Our initial statement of Information-Sensitive Value built in a second claim, specifying the decision-theoretic relationship between the value of an action and the value of its consequences:

Value is Expectational: The value of S's X-ing at t is its expected value according to a relevant probability function Pr and utility function V.

Together with the claim that Value is Information-Sensitive, the claim that Value is Expectational yields Information-Sensitive Value by grounding the decision-theoretic claims made by Information-Sensitive Value.

The claim that Value is Expectational is incompatible with most views that allow value incommensurability, since the resulting decision theories cannot be given an expected value representation. Precisely for this reason, the claim that Value is Expectational must be replaced by an alternative decision theory supplied by the incommensurability theorist. But this is not particularly surprising, because to ask the maximizing theorist to make room for incommensurability just is to ask her to make room for alternatives to expected value theory. The maximizing theorist may not think that we need to do this, but she can make room by doing precisely as requested, removing the claim that Value is Expectational and retaining the rest of her theory.

In this section, we considered how maximizing consequentialists should react to cases in which there is no best option, because all options are dominated. We saw that they might plausibly treat the cases as rational dilemmas by appealing to a minimal dominance principle. Alternatively, consequentialists could treat these cases as a rare exception to maximizing normative theories. That would be disappointing, perhaps, but not a substantial change or the sort of thing to scuttle the broader theory. In the next section, I discuss a significant change that has been proposed to maximizing consequentialism. I argue that traditional philosophical ways of motivating this change are orthogonal to discussions of bounded rationality.

4 Orthogonal challenges

Suppose you are at the grocery store choosing among a shelf of artisanal vinegars. Making a satisfactory choice is easy: just choose any bottle that is in-budget and has a few reliable

indicators of quality. But making an optimal choice is very hard. Are you really required to choose the best, or expectedly best bottle of vinegar?

Michael Slote (1984; 2004) takes the moral of such cases to be that the Maximizing Criterion of Correctness built into maximizing consequentialism should be replaced with a weaker satisficing criterion:

> **Satisficing Criterion of Correctness:** S may (must) X at t just in case X is sufficiently good (the only sufficiently good option) for S at t.

Slote took himself to be motivated by Simon's work on bounded rationality, but on reflection it is not clear that Slote's proposal or the bulk of the ensuing philosophical debate (Byron 2004; Tucker 2016; Yetter Chappell 2019) turns strongly on considerations of bounded rationality.

Consider one of the cases that Slote (1984) uses to motivate his Satisficing Criterion of Correctness. In this case, a home-seller sets the price of the home at a satisfactorily high number. When asked why she didn't ask for more, she replies that she is a satisficer: for her, enough is enough. Importantly, Slote asks us to imagine that the satisficer is not concerned with the probability of making a sale, the time and effort involved in making a sale, or other traditional mainstays of bounded rationality. Slote's satisficer is not concerned that she could not get more money, or that it would be costly to do so, but simply has no interest in receiving more money. The Satisficing Criterion of Correctness is meant to explain why it is permissible for Slote's satisficer to do this.

Bounded rationality theorists can respond to such cases in a variety of ways. Most directly, they might deny that it is rational for Slote's satisficer to turn down free money (Dreier 2004; Mulgan 2001; Thorstad 2022). On this view, it is better to have more money than less, if nothing else because the extra money can be given to charity. But more to the point, they should note that Slote's case is explicitly and deliberately constructed to rule out most leading considerations of bounded rationality. Slote directly stipulates that the agent is not concerned with time, effort, or other bounds. This suggests that Slote aimed

to capture a motivation for satisficing distinct from bounded rationality. What else might Slote have been after?

Some readers have thought that Slote was making a point about the virtue of moderation: a moderate individual would have moderate desires, and hence would not want excessive amounts of money (Swanton 1993, 2004). Indeed, Slote (2001) himself later suggests this interpretation, linking moderation to a species of Stoic virtue. This claim would not pose any threat to the Maximizing Criterion of Correctness, which is a claim about rightness and not a claim about virtue. This claim would only threaten the Maximizing Criterion of Correctness on views which used claims about virtue to ground claims about right action (Hursthouse 2001). Even then the threat is not so clear: absent a strong unity of virtues thesis (Penner 1973), from the fact that an action would exhibit one virtue, such as moderation, it does not follow that a virtuous person would do it.

Others have taken Slote to be making a point about the demandingness of consequentialism: it is wrong to demand too much sacrifice of agents. For example, Slote holds that it is permissible for a hotelier to give a free room to a poor family, but not to give them the best room. Limiting the demandingness of consequentialism is a live challenge to maximizing consequentialism which runs deep through philosophical debates about satisficing (Goodin 2008; Yetter Chappell 2019). But it is not clear that this challenge turns on considerations of bounded rationality: it is fully possible to defend theories of bounded rationality on which morality is maximally demanding, or theories of unbounded rationality on which morality is less demanding, and indeed many theorists have done both.

Perhaps the gap between bounded rationality and satisficing consequentialism is clearest on views such as Richard Chappell's (2019) willpower satisficing. This holds, relative to an effort ceiling X, that:

> **Willpower-Satisficing Consequentialism** An act A is permissible iff it produces no less utility than any alternative action the agent could perform at the cost of up to either (i) X willpower, or (ii) the willpower cost of A, whichever

is greater. (Yetter Chappell 2019, p. 255)

Whatever the fate of such views, we should not take them to be aimed at capturing the full demands of bounded rationality. Otherwise, we would be saddled with the implausible consequence that limitations on willpower are the only bounds that matter rationally, whereas limitations on computation, attention, memory and the like do not. Such a view stands in stark contrast to how the bounded rationality project is traditionally understood.

In this section, we have seen that a traditional way of motivating a Satisficing Criterion of Correctness does pose a challenge to maximizing consequentialism, but that this challenge is largely orthogonal to debates about bounded rationality. It is best understood as a debate about something else, such as the demandingness of morality or the virtue of moderation. The next section considers how maximizing views make sense of the rationality of heuristic cognition.

5 Level separation

One of the most widely accepted claims about bounded rationality is that it is often rational for agents to make decisions using a toolbox of fast-and-frugal heuristics (Gigerenzer and Selten 2001; Gigerenzer and Gaissmaier 2011). Heuristics do not involve explicitly calculating expected values. Heuristics often make decisions based on a small subset of the available evidence, or draw only some of the many inferences warranted by that evidence. But maximizing consequentialism seems to require agents to explicitly calculate the expected values of all available options before acting. If that is true, then maximizing consequentialism is in tension with the rationality of heuristic cognition. That would be a bad consequence. How can maximizing consequentialism be squared with the rationality of heuristic cognition?

5.1 Level separation

To see the way forward, consider a traditional complaint against maximizing consequentialism. Maximizing consequentialism seems to require agents to explicitly calculate the expected utility of every option they face. If this is right, then it raises two objections.

The first is the *computational demandingness objection* (Williams 1973a). Explicitly calculating expected utilities is often expensive, and sometimes entirely beyond an agent's ability. When calculating expected utilities is more trouble than it is worth, maximizing consequentialism seems to demand cognitive wastefulness. And when calculating expected utilities goes beyond an agent's abilities, maximizing consequentialism threatens to violate the principle that ought implies can.

A related objection is the *alienation objection* (Stocker 1976; Williams 1973b). Agents who always make decisions by coldly calculating expected utilities will alienate those around them. Their families, friends and kin will feel unloved. Society at large will distrust them. For these reasons, it would often be better for agents to make decisions without calculating expected utilities, perhaps allowing themselves to be moved to act directly out of love. But if that is right, then consequentialism appears self-defeating, since maximizing consequentialism says to always do what is best.

The traditional response to these worries is to posit a sharp *level separation* (Kagan 2000; Parfit 1984; Thorstad 2024) between two questions: the lower-order question of which nondeliberative acts an agent ought to take, and the higher-order question of how she ought to deliberate about which nondeliberative act to take. For example, it is one question whether I ought to buy coffee and another question entirely how I ought to deliberate about whether to buy coffee. We can apply maximizing consequentialism directly to both questions: I ought to buy coffee just in case it is expectedly best to buy coffee, and I ought to deliberate using the process which it would be expectedly best for me to use.

Level separation solves the computational demandingness objection by letting costs and abilities bear directly on the question of how I ought to deliberate. If it is too costly

to explicitly calculate the expected utility of buying coffee, then precisely for that reason I should use a simpler decision procedure. And if I cannot explicitly calculate the expected utility of buying coffee, then explicitly calculating expected utilities is not an option for me, and hence I cannot be required to do it.

Level separation solves the alienation objection by letting alienation bear on the rightness of decision procedures. If the decision procedure of coldly calculating decision processes would breed hurt and distrust, then that is a strong reason not to use it. It might be better to decide by allowing myself to be moved directly out of love. In this case, maximizing consequentialism says that I should decide in this way.

One feature of level separation is that it allows for mismatch between the nondeliberative actions I ought to take and the outputs of the deliberative processes I ought to use. It might be that I ought to buy coffee, because it is expectedly best to buy coffee. It might also be that I ought to deliberate about whether to buy coffee using some heuristic H, because H is the expectedly best decision process given a full specification of its speed, cost, reliability and other factors. And it might be that although H is generally reliable, H will, alas, misfire in this case, and output the decision not to buy coffee.

This type of *level tension* has long been noted as a feature of the orthodox consequentialist view (Kagan 2000; Parfit 1984; Thorstad 2024). It may make some theorists uncomfortable, but it has been defended at length, and some theorists have gone further, claiming that the ability to make sense of level tension is a virtue of the consequentialist view (Driver 2001; Thorstad 2024). Because my interest in this paper is not to throw sparks on old debates, but rather to evaluate the novel contributions of bounded rationality, I want to take the orthodox consequentialist solution on board and ask how that solution, if sound, could make sense of the rationality of heuristic cognition.

5.2 Applying level separation

A traditional objection to maximizing theories of bounded rationality is that they cannot accommodate the rationality of heuristic cognition (Gigerenzer and Selten 2001b). In-

deed, in one traditional usage, satisficing denotes a particular decisionmaking heuristic (Gigerenzer and Gaissmaier 2011; Selten 1998; Simon 1955, 1956). Suppose you are deciding between destinations for your next vacation. Satisficers begin by fixing an *aspiration level* in one or more goods, which options will be required to meet. Perhaps a destination must be within 500 miles of your home, cost less than $500 per night, and have a zoo. Satisficers then search for options one at a time, comparing each option to the aspiration level. One an option, in this case a destination, is found which meets all of the agent's aspirations, she halts decisionmaking and chooses that option.

A range of theoretical (Dawes and Corrigan 1974; Martignon 2001) and experimental (Gigerenzer et al. 1991; Goldstein and Gigerenzer 2002) results demonstrates that heuristic decision processes often perform very well, returning high-quality decisions at low cognitive cost. At least three reasons are given for thinking that heuristic decision processes are often rational. First, in many cases there is an *accuracy-effort tradeoff* between the quality of decisions and the effort of making them (Johnson and Payne 1985). Heuristics often perform well on the accuracy-effort tradeoff, returning high-quality decisions at low cognitive cost. Second, sometimes it is beyond an agent's ability to use complex nonheuristic methods. In this case, heuristics are the only game in town. Finally, in some contexts heuristics reliably outperform nonheuristic methods, even ignoring cognitive costs, due to *less-is-more effects* (Geman et al. 1992; Gigerenzer and Brighton 2009). In such cases, heuristics are better in every way than nonheuristic methods, returning better decisions at lower cost.

Maximizing consequentialism accommodates the rationality of heuristic cognition by grounding all three justifications for the rationality of heuristics. Maximizing consequentialism incorporates accuracy-effort tradeoffs because cognitive resources wasted on one inquiry could have been better spent on another. Maximizing consequentialism incorporates limited cognitive abilities, because ought implies can. If it is not within our ability to use some cognitive process, then maximizing consequentialism does not require us to do so. And maximizing consequentialism incorporates less-is-more effects, because in the

presence of less-is-more effects, heuristics have better expected consequences of all types than nonheuristic methods.

But how can heuristics be rational, if they sometimes output decisions that do not maximize expected utility? Here the consequentialist relies on level separation. It can be rational to use heuristic decision processes, even though those processes sometimes lead to suboptimal actions. In such cases, maximizing theories say that the agent used a rational decision procedure, but was led to take an irrational action as a result. We saw in Section 5.1 that such cases involve a familiar type of level separation which has been defended at length by consequentialists. It may be that some bounded rationality theorists do not like to admit level tension into normative theorizing, but unless we are provided with a novel argument against level tension, it is not clear that bounded rationality adds anything new to the debate.

Level separation allows us to make sense of a range of related challenges to maximizing consequentialism. For example, it is held that requiring agents to use the expectedly best decision process would make an agent's choice problem too complex, requiring them agents to explicitly calculate the expected values of each decision process (Elster 1983; Lin 2014; Selten 2001). But this is not what maximizing consequentialism says. Maximizing consequentialism posits another level separation between the second-order question of which deliberative process agents should use, and the third-order or metacognitive question of how they should decide which deliberative process to use (Thorstad 2022). It is quite right that the metacognitive process of explicitly calculating expected utilities of each lower-order decision process is often more trouble than it is worth. Precisely for this reason, maximizing consequentialism says that agents should often use simpler metacognitive decision processes. For example, many authors think that agents can often make adequate metacognitive choices by allowing themselves to be directly guided by a special class of noetic feelings (Proust 2013) without any explicit calculation at all. When this is a frugal and reliable strategy, maximizing consequentialism may recommend it.

In this section, we have seen how a traditional level separation between the rationality

of processes and the rationality of their outputs recovers three traditional arguments for the rationality of heuristic cognition and answers concerns about the complexity of consequentialist metacognition. In the next section, I discuss the argument that maximizing standards are inappropriate under Knightian uncertainty.

6 Knightian uncertainty

One of the best known arguments against maximization for bounded agents is that under conditions of deep or 'Knightian' uncertainty, maximizing standards are inappropriate (Gigerenzer and Selten 2001b; Klein 2001). Knightian uncertainty is a difficult issue to broach, for two reasons. First, many things have been meant by Knightian uncertainty, so it is important to keep track of what we are talking about. And second, this is an area of theorizing in which it is often difficult to move established positions through argument. Nevertheless, in this section I explore two ways of filling out the notion of Knightian uncertainty that could underlie an objection to maximization. I argue that plausible ways of developing each conception of Knightian uncertainty leave substantial room for bounded agents to be subject to maximizing normative standards.

6.1 Unawareness

Many bounded rationality theorists understand Knightian uncertainty to involve unawareness of possible states that the world could be in (Artinger et al. 2014; Mousavi and Gigerenzer 2014; Gigerenzer and Sturm 2012). They hold that maximization is impossible or inappropriate under conditions of unawareness. But precisely what is the problem with maximizing under conditions of unawareness?

Sometimes the objection is simply that heuristics are uniquely appropriate under conditions of unawareness (Mousavi and Gigerenzer 2017). But this is no objection to maximizing consequentialism. We saw in Section 5 that maximizing consequentialism is fully compatible with the rationality of heuristic cognition.

Perhaps the objection is that agents who maximize expected utility according to less-aware probability and utility functions will often fail to maximize expected utility relative to a fully-aware probability and utility function. But at least from the *ex ante* perspective, this is a feature of the view, not a bug. Bounded agents are ignorant in two ways: they lack information about how the world is, and also about how they world could be.[5] Maximizing consequentialism incorporates agents' limited information about how the world is by calculating expected utilities relative to an agent's evidence about the world. Similarly, the natural extension of maximizing consequentialism incorporates agents' limited information about how the world could be by calculating expected utilities relative to the world states that they are aware of.[6] Theories of bounded rationality would be wrongly insensitive to an agent's limited information about how the world is if they did not restrict attention to expected utilities. And likewise, *ex ante* theories of bounded rationality would by wrongly insensitive to an agent's limited information about how the world could be if they did not restrict attention to states of which agents are relevantly aware.

Sometimes the objection is made that under conditions of unawareness, the notions of probability and utility become ill-defined (Artinger et al. 2014; Mousavi and Gigerenzer 2014; Gigerenzer and Selten 2001b). But this objection requires separate treatment. The problem cannot be that probability and utility functions are formally undefinable under conditions of unawareness. Recent work in Bayesian decision theory has shown how probability and utility functions can be defined relative to a restricted state of awareness and how they might change with growing awareness (Bradley 2017; Karni and Vierø 2013; Steele and Stefánsson 2021). Furthermore, this objection is driven by many concerns beyond unawareness. So it seems most prudent to group this objection with a broader

[5]There are some technical notions of information on which unawareness is not a type of informational lack. It is not my aim to argue over terminology: those who hold this view should replace 'information about' with 'awareness of' in relevant parts of the subsequent discussion.

[6]This is not quite true. For example, on a popular 'catchall' model, agents also assign probabilities and utilities to a 'catchall' state capturing all possibilities they are not aware of. But this does not change the general point that maximizing consequentialism, as I have formulated it, captures a broadly *ex ante* perspective in which agents calculate expected values based on limited information and awareness.

view of Knightian uncertainty that involves the ill-definedness of probabilities or utilities.

6.2 Undefinedness

Some bounded rationality theorists hold that under Knightian uncertainty, probabilities or utilities become undefined (Artinger et al. 2014; Mousavi and Gigerenzer 2014; Gigerenzer and Selten 2001b). To motivate this view, consider:

> **(Jellyfish)** A stranger approaches you and begins pulling objects from a sack. He takes out a tube of toothpaste, a jellyfish, and then a smaller tube of toothpaste. How likely is it that the stranger will take out another tube of toothpaste? (Elga 2010)

It may seem compelling to think that there is no fact about how likely it is that the stranger will take out another tube of toothpaste. Given the depth of our uncertainty about what the stranger will do, it might appear that the relevant probabilities are so unclear that they should be treated as ill-defined.

Three observations may serve to limit the class of cases on which this view gets a take. First, note that the ill-definedness of probabilities and utilities poses a radical challenge to a great number of important philosophical notions. For example, the notion of evidential probability used by maximizing consequentialism is also a mainstay of traditional epistemology. And the notion of utility used by maximizing consequentialism is the same notion that makes axiological theorizing possible. To deny that evidential probabilities or utilities are defined in some case is thereby to deny that important philosophical projects get a take on that case. While this is a consequence we might be willing to tolerate in some rare cases, it is not a consequence to be taken lightly.

Second, some ways of treating probabilities and utilities as ill-defined may retain a substantial core of maximizing consequentialism. For example, we might take an imprecise probabilistic approach on which single probability and utility functions are replaced with sets of probability and utility functions (Troffaes 2007; Walley 1991). While some

decision rules for imprecise probabilities remove the requirement to maximize expected utility, others retain analogs of this requirement. For example, we could require agents to take acts which maximize expected utility on some or every probability and utility function in the relevant set.[7] That would be a change to the statement of maximizing consequentialism in Section 2, but it would not be so large a change as we might have expected.

Finally, in specifying the scope of Knightian uncertainty, the objection from undefinedness faces a dilemma. To illustrate the problem, consider a famous taxonomy of decision situations due to Luce and Raiffa (Luce and Raiffa 1957, p. 13). On this taxonomy, *decisionmaking under certainty* occurs when the consequences of all acts are known; *decisionmaking under risk* occurs when some consequences are unknown but their probabilities are known; and *decisionmaking under uncertainty* occurs when outcome probabilities are "completely unknown, or not even meaningful". Where, in this taxonomy, should we put Knightian uncertainty?

Taking the first horn of the dilemma involves giving a broad conception of Knightian uncertainty. For example, we could say that Knightian uncertainty involves anything deeper than decisionmaking under risk. This means that even when probabilities are based on broad and robust evidence, unless every one of the relevant probabilities meets the demanding standards to constitute knowledge, we face Knightian uncertainty.[8] But once we take such a broad conception of Knightian uncertainty, the idea that probabilities or utilities become undefined under Knightian uncertainty loses its appeal. Most philosophers are willing to accept that evidential probabilities are well-defined in cases where an agent's evidence base is broad and robust. If the objection from Knightian uncertainty involves denying this claim, then we are well within our rights to ask for further argument.

Taking the second horn of the dilemma involves giving a narrow conception of Knightian uncertainty. At the extreme, we could identify Knightian uncertainty with Luce and

[7]See Troffaes (2007) for a survey of these and other common decision rules.

[8]Even if we follow Moss (2018) in thinking that credences can constitute knowledge, we needn't think that all decision-relevant credences typically do constitute knowledge.

Raiffa's conditions of decisionmaking under uncertainty, under which probabilities are completely unknown, or not even meaningful. Or, broadening a little, we might take Knightian uncertainty to obtain in conditions of especially deep uncertainty such as Jellyfish, and also in problems such as long-run climate modeling (Lempert et al. 1996; Ranger et al. 2013). Because these conceptions of Knightian uncertainty are highly restricted, they lend more plausibility to the idea that probabilities and utilities could be ill-defined under Knightian uncertainty. But precisely because these conceptions of Knightian uncertainty are so narrow, they leave open a great number of cases in which bounded agents do not face Knightian uncertainty. And in those cases, there will be no objection to the applicability of maximizing normative standards.

Summing up, in this section we considered the objection that maximizing is inappropriate under conditions of Knightian uncertainty. We looked at two ways in which the notion of Knightian uncertainty can be understood: as unawareness, or as ill-definedness of probabilities and utilities. We saw that in both cases, there is no decisive reason given to relax maximizing normative standards in a large number of interesting cases of rational decisionmaking for bounded agents. As a result, current appeals to Knightian uncertainty cannot sustain the idea that maximizing is always, or even usually inappropriate for bounded agents.

7 Philosophical implications

This paper argued that the rationality of satisficing can be squared with a substantial core of traditional maximizing normative theories. By way of illustration, we considered the extent to which traditional arguments for satisficing could be reconciled with maximizing consequentialism.

Maximizing Consequentialism: For all agents S, acts X and times t:

> **Maximizing Criterion of Correctness:** S may (must) X at t just in case X is maximal (uniquely maximal) for S at t.

> **Information-Sensitive Value:** The value of S's X-ing at t is the expected value of X's consequences, where the expectation in question is taken relative to evidential probabilities. That is, $V(X) = \sum_w Pr_E(w)V(X,w)$, where Pr_E are the evidential probabilities generated by S's total evidence at t.

We saw in Section 3 that some objections, based on cases in which all options are dominated or some values are incommensurable, occasion at most modest and well-motivated revisions to Maximizing Consequentialism. We saw in Section 4 that traditional philosophical arguments for a satisficing criterion of correctness draw largely on concerns such as a desire to limit the demands of morality or to elucidate the virtue of moderation. These concerns, while philosophically important, do not turn essentially on considerations of bounded rationality, and as such can be assessed separately. We saw in Section 5 that maximizing theories can incorporate the rationality of decisionmaking heuristics such as satisficing by invoking a traditional level separation between the rationality of decisionmaking processes and the rationality of the intentions or actions they produce. And we saw in Section 6 that on leading conceptions of Knightian uncertainty as unawareness or undefinedness, maximizing criteria may be defended except possibly in the deepest cases of uncertainty. As with cases in which all options are dominated (Section 3.1), such cases are rare and the rational reaction to such cases is disputed, but precisely because they are rare we could accept them as limitations on the scope of maximizing theories without substantial loss.

This discussion suggests that the truth of maximizing theories, such as Maximizing Consequentialism, may be entirely or largely compatible with considerations of bounded rationality thought to rule them out. If that is right, then it has at least four important philosophical implications.

7.1 A neglected normative alternative

Many philosophical theories of bounded rationality take as their point of departure the idea that traditional maximizing standards must be relaxed to make room for bounded rationality. For example, say that an act is *apparently appropriate* if it maximizes expected utility given an agent's beliefs and desires. Instead of requiring agents to take all apparently appropriate actions:

> **(Full Rationality)** If A has a particular belief-desire set, A would undertake *all* and only actions that are apparently appropriate. (Cherniak 1986, p. 8)

we might only require agents to take some apparently appropriate actions:

> **(Minimal Rationality)** If A has a particular belief-desire set, A would undertake some, but not necessarily all, of those actions that are apparently appropriate. (Cherniak 1986, p. 8)

Or instead of a maximizing criterion of correctness (Section 2) we might impose a satisficing criterion of correctness (Slote 1984; Yetter Chappell 2019):

> **Satisficing Criterion of Correctness:** S may (must) X at t just in case X is sufficiently good (the only sufficiently good option) for S at t.

Programs such as Slote's satisficing consequentialism and Cherniak's minimal rationality exert a continued influence on normative theorizing across a number of fields.[9]

What has received less attention is the possibility of taking a maximizing approach to the study of bounded rationality. If this paper is on the right track, then maximizing approaches deserve a second look. And if I am right to claim that considerations of bounded rationality do not force us to significantly revise maximizing normative standards, then

[9]For example, minimal rationality has recently been applied to the epistemology of inquiry (Hoek forthcoming) and belief (Ganapini 2020), and satisficing approaches have been explored in moral philosophy (Yetter Chappell 2019) and the philosophy of religion (Tucker 2016, 2017).

we cannot infer the need for programs such as minimal rationality and satisficing consequentialism on the basis of bounded rationality alone. Those programs will need to be motivated on other grounds.

7.2 Bayesianism: new hope, and new challenges

One of the fastest-growing paradigms in cognitive science is the Bayesian approach. Fueled by new mathematical tools and theoretical accounts, Bayesian tools are enjoying a renaissance across fields as diverse as perception (Kersten et al. 2004), reasoning (Elqayam and Over 2013) and memory (Anderson 1990).

Because Bayesian approaches are built around a maximizing normative theory, opponents of maximizing may have expected Bayesian views to deliver inadequate accounts of the behavior of bounded agents. Early Bayesians did struggle to account for bounded rationality, because they neglected to build relevant bounds into their models. For example, a paradigmatic early statement of the Bayesian method of rational analysis comes with the injunction to make minimal assumptions about computational limitations (Anderson 1990; Chater and Oaksford 1999). This is surprising. Agents have computational limitations, and those computational limitations bear on how they can and should cognize. Couldn't we build better descriptive and normative models by incorporating computational limitations directly into our models?

Recent Bayesian paradigms such as cognitively bounded rational analysis (Howes et al. 2009), boundedly rational analysis (Icard 2018) and resource-rational analysis (Lieder and Griffiths 2020) have learned this lesson, building computational limitations and other cognitive bounds directly into Bayesian models. As a result, Bayesian models are increasingly able to deliver compelling descriptive and normative claims about bounded rationality. If the argument in this paper is on the right track, these newer Bayesian models should be taken seriously as accounts of bounded rationality.

However, recent Bayesian accounts have succeeded by using new explanatory strategies that have important differences from traditional Bayesian accounts. For example,

agents might make decisions based on samples of information retrieved from memory, rather than by exhaustively surveying all of their evidence (Stewart et al. 2006; Vul et al. 2014). Agents could use heuristics to approximate Bayesian calculations (Oaksford and Chater 2007). And agents might respond in a utility-maximizing way to the costs of cognition (Lieder and Griffiths 2020). These changes will have consequences that may be unfamiliar territory for traditionally-minded Bayesians. For example, agents who cognize heuristically will be vulnerable to forming incoherent beliefs, and agents who make decisions based on samples of evidence will sometimes judge against the balance of their total evidence. In this way, while taking maximization seriously opens the possibility for taking seriously Bayesian accounts of bounded rationality, successful Bayesian accounts of bounded rationality may be importantly different from their predecessors.

Finally, the discussion in this paper contributes to recent calls to clarify the explanatory status of Bayesian principles (Colombo and Hartmann 2017; Colombo et al. 2021; Levy 2023). The thesis I have defended in this paper is normative, not descriptive. I said that maximizing consequentialism may characterize how bounded agents ought to behave. Bayesian cognitive science must go beyond this bare normative proposal, since cognitive scientists have descriptive ambitions. Although recent Bayesians have attempted to say how their models give lower-level 'algorithmic' explanations of cognitive processes (Griffiths et al. 2015), some critics worry that the explanatory status of Bayesian models reflects their ability to specify the higher-order computational task that cognitive processes aim to perform (Bowers and Davis 2012; Jones and Love 2011), and even leading Bayesians concede that some early models provided primarily computational explanations (Tenenbaum et al. 2011). It is of the utmost importance to get clear on how, if at all, this talk of computational-level explanation differs from specifying the normative standard that agents aim to meet, otherwise some leading Bayesian accounts will fail to advance substantially beyond the normative thesis defended in this paper.

7.3 Maximizing what?

In this paper, I argued that a particular maximizing normative theory, namely maximizing consequentialism, can account for insights from bounded rationality without substantial change. But many theorists are not consequentialists. For example, some philosophers (Stich 1990) have defended instrumentalist theories of bounded rationality on which agents should maximize their own interests, but not the interests of others. I suspect that many of the arguments in this paper would also succeed in rescuing instrumentalist accounts.

However, other theorists defend different accounts. For example, Julia Staffel (Staffel and De Bona 2018; Staffel 2020, 2023) and perhaps also Lyle Zynda (Zynda 1996) defend a coherence-based approach on which agents should strive to be as coherent as possible, given their bounds. Bounded rationality theorists have sometimes expressed skepticism about the explanatory prospects of coherence-based approaches (Arkes 2016; Thorstad forthcoming a,b). It is worth thinking carefully about whether, and to what extent the arguments in this paper might extend to make room for coherence-based approaches to bounded rationality.

7.4 The meaning of 'satisficing'

The term satisficing is used in many ways. Sometimes, it is used to denote a particular decisionmaking heuristic, as we saw in Section 5. Other times, it is used to denote a novel standard of correctness, as we saw in Section 4. And these may not exhaust the uses of satisficing. For example, we saw in Section 4 that satisficing may be used to describe a pattern of moderate preference. The same usage reappears in studies of consumer choice, where it is claimed that satisficers, who are content with good products and not obsessed with finding the best, make better choices by avoiding decision paralysis when faced with large menus of options (Scheibehenne et al. 2010).

It is of the utmost importance to clarify and distinguish different uses of satisficing,

because some versions of satisficing will be consistent with traditional maximizing theories whereas others will not. For example, we saw that satisficing criteria of correctness conflict with maximizing criteria, but that neither the rationality of satisficing as a decision heuristic nor the claim that satisficing typifies a virtuous pattern of moderate preference conflicts with traditional maximizing theories. Likewise, when consumer choice theorists claim that agents who place little emphasis on finding the best products make better choices, this could be interpreted as a claim that these agents' patterns of preference or desire are rational in crowded choice environments because they lead to quicker and better choices. To the extent that we are willing to pronounce on the rationality of preferences or desires, this claim would be compatible with traditional maximizing theories on much the same grounds given for the rationality of heuristic preference.

I do not claim that this survey of uses of satisficing is exhaustive. Indeed, I am quite sure that it is not. For this reason, one pressing project for future work would be to catalog and distinguish the many various uses of satisficing, then to explore the normative relationships between each claim and traditional philosophical theories.

References

Anderson, John. 1990. *The adaptive character of thought*. Lawrence Erlbaum Associates.

Arkes, Hal, Gigerenzer, Gerd, and Hertwig, Ralph. 2016. "How bad is incoherence?" *Decision* 3:20–39.

Artinger, Florian, Petersen, Malte, Gigerenzer, Gerd, and Weibler, Jürgen. 2014. "Heuristics as adaptive decision strategies in management." *Journal of Organizational Behavior* 36:533–52.

Bermúdez, José. 2009. *Decision theory and rationality*. Oxford University Press.

Bowers, Jeffrey and Davis, Colin. 2012. "Is that what Bayesians believe? Reply to Griffiths, Chater, Norris and Pouget (2012)." *Psychological Bulletin* 138:423–426.

Bradley, Richard. 2017. *Decision theory with a human face*. Cambridge University Press.

Brighton, Henry. 2020. "Rationality without optimality: Bounded and ecological rationality from a Marrian perspective." In Riccardo Viale (ed.), *Routledge handbook of bounded rationality*, 268–79. Routledge.

Byron, Michael (ed.). 2004. *Satisficing and maximizing: Moral theorists on practical reason*. Cambridge University Press.

Chappell, Richard. 2019. "Willpower satisficing." *Nous* 53:251–65.

Charlow, Nate. 2013. "What we know and what to do." *Synthese* 190:2291–323.

Chater, Nick and Oaksford, Mike. 1999. "Ten years of the rational analysis of cognition." *Trends in Cognitive Sciences* 3:57–65.

Cherniak, Christopher. 1981. "Minimal rationality." *Mind* 90:161–183.

—. 1986. *Minimal rationality*. MIT Press.

Colombo, Matteo, Elkin, Lee, and Hartmann, Stephan. 2021. "Being realist about Bayes, and the predictive processing theory of mind." *British Journal for the Philosophy of Science* 72:185–220.

Colombo, Matteo and Hartmann, Stephan. 2017. "Bayesian cognitive science, unification, and explanation." *British Journal for the Philosophy of Science* 68:451–84.

Dawes, Robyn and Corrigan, Bernard. 1974. "Linear models in decision making." *Psychological Bulletin* 81:95–106.

Dreier, James. 2004. "Why ethical satisficing makes sense and rational satisficing doesn't." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 131–154. Cambridge University Press.

Driver, Julia. 2001. *Uneasy Virtue*. Cambridge University Press.

Elga, Adam. 2010. "Subjective probabilities should be sharp." *Philosophers' Imprint* 10:1–11.

Elqayam, Shira and Over, David. 2013. "New paradigm psychology of reasoning: An introduction to the special issue edited by Elqayam, Bonnefon, and Over." *Thinking and Reasoning* 19:249–65.

Elster, Jon. 1983. *Sour grapes: Studies in the subversion of rationality*. Cambridge University Press.

Ganapini, Marianna Bergamaschi. 2020. "Belief's minimal rationality." *Philosophical Studies* 177:3263–3282.

Geman, Stuart, Bienenstock, Elie, and Doursat, René. 1992. "Neural networks and the bias/variance dilemma." *Neural Computation* 4:1–58.

Gigerenzer, Gerd and Brighton, Henry. 2009. "Homo heuristicus: Why biased minds make better inferences." *Topics in Cognitive Science* 1:107–43.

Gigerenzer, Gerd and Gaissmaier, Wolfgang. 2011. "Heuristic decision making." *Annual Review of Psychology* 62:451–82.

Gigerenzer, Gerd and Goldstein, Daniel. 1996. "Reasoning the fast and frugal way: Models of bounded rationality." *Psychological Review* 103:650–9.

Gigerenzer, Gerd, Hoffrage, Ulrich, and Kleinbölting, Heinz. 1991. "Probabilistic mental models: A Brunswikian theory of confidence." *Psychological Review* 98:506–528.

Gigerenzer, Gerd and Selten, Reinhard (eds.). 2001. *Bounded rationality: The adaptive toolbox*. MIT Press.

Gigerenzer, Gerd and Selten, Reinhard. 2001b. "Rethinking rationality." In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality: The adaptive toolbox*, 1–12. MIT Press.

Gigerenzer, Gerd and Sturm, Thomas. 2012. "How (far) can rationality be naturalized?" *Synthese* 187:243–68.

Goldstein, Daniel and Gigerenzer, Gerd. 2002. "Models of ecological rationality: The recognition heuristic." *Psychological Review* 109:75–90.

Goodin, Robert. 2008. "Demandingness as a virtue." *Journal of Ethics* 13:1–13.

Griffiths, Thomas L., Lieder, Falk, and Goodman, Noah. 2015. "Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic." *Topics in Cognitive Science* 7:217–29.

Hajek, Alan. 2014. "Unexpected expectations." *Mind* 123:533–67.

Hoek, Daniel. forthcoming. "Minimal rationality and the web of questions." In Dirk Kindermann, Peter van Elswyk, and Andy Egan (eds.), *Unstructured content*. Oxford University Press.

Howes, Andrew, Lewis, Richard, and Vera, Alonso. 2009. "Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action." *Psychological Review* 116:717–51.

Hursthouse, Rosalind. 2001. *On virtue ethics*. Oxford University Press.

Icard, Thomas. 2018. "Bayes, bounds, and rational analysis." *Philosophy of Science* 85:79–101.

Johnson, Eric and Payne, John. 1985. "Effort and accuracy in choice." *Management Science* 31:395–414.

Jones, Matt and Love, Bradley. 2011. "Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition." *Behavioral and Brain Sciences* 34:169–231.

Kagan, Shelly. 2000. "Evaluative focal points." In Brad Hooker, Elinor Mason, and Dale E. Miller (eds.), *Morality, rules and consequences*, 134–55. Edinburgh University Press.

Karni, Edi and Vierø, Marie-Louise. 2013. "'Reverse Bayesianism': A choice-based theory of growing awareness." *American Economic Review* 103:2790–2810.

Kersten, Daniel, Mamassian, Pascal, and Yuille, Alan. 2004. "Object perception as Bayesian inference." *Annual Review of Psychology* 55:271–304.

Klein, Gary. 2001. "The fiction of optimization." In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality: The adaptive toolbox*, 103–21. MIT Press.

Kolodny, Niko and MacFarlane, John. 2010. "Ifs and oughts." *Journal of Philosophy* 107:115–43.

Lempert, Robert, Schlesinger, Michael, and Bankes, Steve. 1996. "When we don't know the costs or the benefits: Adaptive strategies for abating climate change." *Climatic Change* 33:235–74.

Levy, Arnon. 2023. "Can Bayesian models of cognition show that we are (epistemically) rational?" *Philosophy of Science* 90:1069–78.

Lieder, Falk and Griffiths, Thomas. 2020. "Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources." *Behavioral and Brain Sciences* 43:E1.

Lin, Hanti. 2014. "On the regress problem of deciding how to decide." *Synthese* 191:661–670.

Luce, R. Duncan and Raiffa, Howard. 1957. *Games and decisions: Introduction and critical survey*. Wiley.

Martignon, Laura. 2001. "Comparing fast and frugal heuristics and optimal models." In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality: The adaptive toolbox*. MIT Press.

Moss, Sarah. 2018. *Probabilistic knowledge*. Oxford University Press.

Mousavi, Shabnam and Gigerenzer, Gerd. 2014. "Risk, uncertainty, and heuristics." *Journal of Business Research* 67:1671–1678.

—. 2017. "Heuristics are tools for uncertainty." *Homo Oeconomicus* 34:361–79.

Mulgan, Tim. 2001. "How satisficers get away with murder." *International Journal of Philosophical Studies* 9:41–46.

Oaksford, Mike and Chater, Nick. 2007. *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.

Parfit, Derek. 1984. *Reasons and persons*. Oxford University Press.

Penner, Terry. 1973. "The unity of virtue." *Journal of Philosophy* 82:35–68.

Pollock, John. 1983. "How do you maximize expectation value?" *Noûs* 17:409–21.

Proust, Joëlle. 2013. *The philosophy of metacognition*. Oxford University Press.

Ranger, Nicola, Reeder, Tim, and Lowe, Jason. 2013. "Addressing 'deep' uncertainty over long-term climate in major infrastructure projects: Four innovations of the Thames Estuary 2100 project." *EURO Journal on Decision Processes* 1:233–262.

Samuels, Richard, Stich, Stephen, and Bishop, Michael. 2002. "Ending the rationality wars: How to make disputes about human rationality disappear." In Renee Elio (ed.), *Common sense, reasoning and rationality*, 236–68. Oxford University Press.

Scheibehenne, Benjamin, Greifeneder, Rainer, and Todd, Peter. 2010. "Can there ever be too many options? A meta-analytic review of choice overload." *Journal of Consumer Research* 37:409–425.

Schmidtz, David. 1995. *Rational choice and moral agency*. Princeton University Press.

—. 2004. "Satisficing as a humanly rational strategy." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 30–58. Cambridge University Press.

Selten, Reinhard. 1998. "Aspiration adaptation theory." *Journal of Mathematical Psychology* 42:191–214.

—. 2001. "What is bounded rationality?" In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality*, 13–37. MIT Press.

Simon, Herbert. 1955. "A behavioral model of rational choice." *Quarterly Journal of Economics* 69:99–118.

—. 1956. "Rational choice and the structure of the environment." *Psychological Review* 63:129–138.

—. 1978. "Rationality as process and as product of thought." *American Economic Review* 68:1–19.

Slote, Michael. 1984. "Satisficing consequentialism." *Proceedings of the Aristotelian Society, Supplement* 58:139–63.

—. 2001. "Moderation and satisficing." In Elijah Millgram (ed.), *Varieties of practical reasoning*, 221–36. MIT Press.

—. 2004. "Two views of satisficing." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 14–29. Cambridge University Press.

Sorensen, Roy. 2006. "Originless sin: Rational dilemmas for satisficers." *Philosophical Quarterly* 56:213–23.

Staffel, Julia. 2020. *Unsettled thoughts: A theory of degrees of rationality*. Oxford University Press.

—. 2023. "Bayesian norms and non-ideal agents." In Maria Lasonen-Aarnio and Clayton Littlejohn (eds.), *Routledge handbook of the philosophy of evidence*, 135–47. Routledge.

Staffel, Julia and De Bona, Glauber. 2018. "Why be (approximately) coherent?" *Analysis* 78:405–15.

Steele, Katie and Stefánsson, Orri. 2021. *Beyond uncertainty*. Cambridge University Press.

Stewart, Neil, Chater, Nick, and Brown, Gordon D.A. 2006. "Decision by sampling." *Cognitive Psychology* 53:1–26.

Stich, Stephen. 1990. *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation.* MIT Press.

Stocker, Michael. 1976. "The schizophrenia of modern ethical theories." *Journal of Philosophy* 73:453–66.

Sturm, Thomas. 2012. "The 'rationality wars' in psychology: Where they are and where they could go." *Inquiry* 55:66–81.

Swanton, Christine. 1993. "Satisficing and virtue." *Journal of Philosophy* 90:33–48.

—. 2004. "Satisficing and perfectionism in virtue ethics." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 176–189. Cambridge University Press.

Tenenbaum, Joshua, Kemp, Charles, Griffiths, Thomas L., and Goodman, Noah. 2011. "How to grow a mind: Statistics, structure, and abstraction." *Science* 331:1279–85.

Thorstad, David. 2022. "Two paradoxes of bounded rationality." *Philosophers' Imprint* 22:1–16.

—. 2024. *Inquiry under bounds*. Oxford University Press.

—. forthcoming a. "The accuracy-coherence tradeoff in cognition." *British Journal for the Philosophy of Science* forthcoming.

—. forthcoming b. "The complexity-coherence tradeoff in cognition." *Mind* forthcoming.

Troffaes, Matthias C.M. 2007. "Decision making under uncertainty using imprecise probabilities." *International Journal of Approximate Reasoning* 45:17–29.

Tucker, Chris. 2016. "Satisficing and motivated submaximization (in the philosophy of religion)." *Philosophy and Phenomenological Research* 93:127–43.

—. 2017. "How to think about satisficing." *Philosophical Studies* 174:1365–84.

Viale, Riccardo. 2021. "Why bounded rationality?" In Riccardo Viale (ed.), *Routledge handbook of bounded rationality*, 1–54. Routledge.

Vul, Edward, Goodman, Noah, Griffiths, Thomas L., and Tenenbaum, Joshua B. 2014. "One and done? Optimal decisions from very few samples." *Cognitive Science* 38:599–637.

Walley, Peter. 1991. *Statistical reasoning with imprecise probabilities*. Chapman and Hall.

Williams, Bernard. 1973a. "A critique of utilitarianism." In J.J.C. Smart and Bernard Williams (eds.), *Utilitarianism: For ane against*, 75–150. Cambridge University Press.

—. 1973b. *Problems of the self*. Cambridge University Press.

Yalcin, Seth. 2016. "Modalities of normativity." In Nate Charlow and Matthew Chrisman (eds.), *Deontic modality*, 230–55. Oxford University Press.

Yetter Chappell, Richard. 2019. "Willpower satisficing." *Noûs* 53:251–65.

Zynda, Lyle. 1996. "Coherence as an ideal of rationality." *Synthese* 109:175–216.